User Experience on USQCD Facilities

Kostas Orginos

The College of William and Mary/Jlab

Overview

User experience on USQCD facilities:

- Obtaining and account
- Compiling codes
  - Hardware targets
- Preparing jobs for submission
- Job monitoring
- Allocation usage
- Data storage
- Problem reporting
- User support from system managers
- User experience
- Conclusions

The QCDOC project at BNL

LQCD Home

QDCOC Computing

Lattice Archives at BNL

Related Publications

Conferences/Meetings

Contacts

LQCD Links

The Lattice Web

US Lattice Gauge Theory

Lattice at Fermilab

Lattice at Jefferson Lab

Columbia University

UKQCD

Visiting BNL

Visitor Information

Lodging near BNL

Directions to BNL

Other Information

QCDOC (Quantum Chromo-Dynamics On-a-Chip) is an international collaboration that was formed to design a massively parallel supercomputer architecture tailored to the needs of of Lattice Quantum Chromo-Dynamics (LQCD) simulations. The project has been centered at Columbia University with contributions from the UKQCD collaboration, the RIKEN BNL Research Center (RBRC), the IBM Thomas J. Watson Research Center and Brookhaven National Laboratory.

Brookhaven National Laboratory (BNL) currently hosts two large (12,288 nodes, 10 TFlops each) QCDOC machines: one for the RBRC community funded by RIKEN (Japan's Institute of Physical and Chemical Research) and the other for the US Lattice Gauge Theory community funded by the US Department of Energy. A third large QCDOC machine for the UKQCD community, funded by PPARC (UK's Particle Physics and Astronomy Research Council), is hosted in Edinburgh, Scotland.

The QCDOC architecture has been designed to provide a highly cost-effective, massively parallel computer capable of focusing significant computing resources on relatively small but extremely demanding problems. The individual processing nodes are PowerPC-based and interconnected in a six-dimensional, low-latency mesh network with the topology of a torus. Each node, designed by our collaboration and built by IBM, includes a single custom ASIC plus DDR SDRAM. It has a peak speed of 1 Gigaflops. More information about this architecture can be found on the QCDOC architecture and publication web pages.

The QCDOC design is a natural evolution of that used in our earlier QCDSP machines (Quantum Chromodynamics on Digital Signal Processors). QCDSP incorporated a low-latency four-dimensional mesh network to realize peak speeds of 1 Teraflops with 20,000 nodes. QCDSP won the Gordon Bell prize at Supercomputing 98 and was acknowledged as the world's fastest non-commercial supercomputer. A separate evolution of the regular mesh architecture of QCDSP is represented by IBM's Blue Gene/L supercomputer, which uses a three-dimensional mesh, will incorporate up to 64,000 processing nodes and has a peak speed of 360 Teraflops.

Fermi National Accelerator Laboratory

Lattice Gauge Theory Computational Facility

Fermilab operates large clusters of computers for lattice quantum chromodynamics, as part of the national computational infrastructure for lattice QCD established by the Department of Energy. Their goal is the understanding of the strong dynamics of quarks and gluons, which is beyond the reach of the traditional perturbative methods of quantum field theory. A central goal of the groups using the computers is the accomplishment of the calculations required to extract from experiment the fundamental parameters of the Standard Model of particle physics.

LQCD Cluster Status

Sun May 6 21:51:48 CDT 2007

PBS : Active Maui : Active dCache : Up 34.15 MB/s (10/240) Free=4627 GB 22%

Physics Program

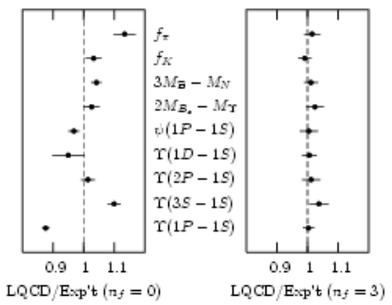

The QCD clusters

The QCD cluster future

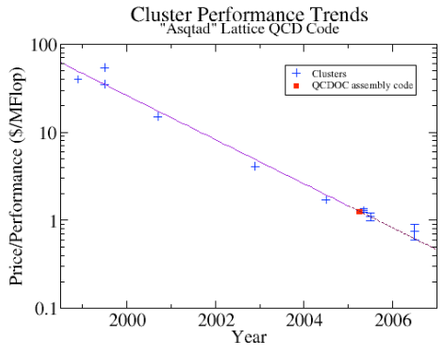

Fermilab Lattice QCD

LQCD Collaboration 2007 All Hands Meeting

Fermilab Cluster Details

DOE, SciDAC Support for Lattice QCD Computing

LQCD User Information:

- New accounts and renewed accounts
- Basic Computer Security Training Requirement
- Notes for Users
- Kerberos Notes for Users
- Additional Kerberos Information
- Transferring Files Between USQCD Sites
- LQCD-Users Mail Archive
- Draft Run Time Environment Specification
- Presentation on User Environment
- Run Time Environment Primer
- RAID Disk Usage

System status

- Cluster Status
- New Muon Temperature
- Help: lqcd-admin@fnal.gov
- LQCD-Admin Mail Archive

Useful code

Cluster work at Fermilab

Lattice links

webmaster@lqcd.fnal.gov

Status

Cluster Queues

mesh ib

Utilization

Job View FS View

Maui Fairshare Stat

Any 06-07 05-06

Project Accounts

Cache Disks

Servers

Computer Room

1st Floor 2nd Floor

News Archive

Documentation

User Environment

USQCD Standard

JLab Specific

Software

QMP QDP++

Chroma

Facilities

Clusters

Disk Cache

Links

Contacts

Jefferson Lab LQCD Homepage

Welcome to the Jefferson Lab Lattice QCD user facilty.
New users start here (last updated March 9th, 2007 now on JLAB's wiki site)

Report a problem

News

30-Apr-07 **Temporary Down Completed --** We are temporarily still constraining usage on the new fileserver in case we in the end need to move back off of it.

30-Apr-07 **qcdi02 rebooted at 3:30 pm --** Rebooted to maintain file system consistency.

25-Apr-07 **/cache/users is available now --** We are moving /cache/users from hpcdata5 to hpcdata6. This work is done at ~9:25am. Now this disk pool can be accessed from all interactive nodes and rcp from all computing nodes.

01-Mar-07 **New debug queues created** -- High priority 30 minute queues are available for the mesh and ib clusters to use them read instructions here.

26-Jan-07 **New problem reporting ticket system available--** A new ticket system to report problems has been implemented. The 'Report a problem' link above makes this system available.

System Status

Menu links to the left can be used to view the queue status.

Mesh Clusters
(512 nodes)

75%

qcdpbs-mesh

Last updated:
May 06 23:58 EDT 2007

Infiniband Clusters
(~320 nodes)

99%

qcdpbs-ib

Last updated:
May 06 23:58 EDT 2007

Feedback (suggestions for new material, links to be added, etc.).

Report a problem

Obtaining an Account

- All three centers have a relative simple procedure
  - Have personally experienced JLAB and FNAL: Worked very well
    - Takes at most a week to get it done (most times faster)
- Web documentation guides the user to the appropriate actions
  - Fill out forms and email requests to the site managers
- National Lab security regulations complicates the process a little

Compilation

- Detailed compilation instructions exist on the web. Pointers to the appropriate hardware related libraries exist.

- Pre-compiled SciDAC software (QIO QLA QDP QDP++) exists (JLAB FNAL BNL).

- Even application code such as Chroma is in some cases pre-compiled

- In certain cases the user just needs to bring in data and input files, write the job scripts and run his jobs. Code compilation already done by the site management (ex. Chroma at JLAB and FNAL).

- Locations of the libraries vary fro site to site. Users would benefit from a uniform pattern.

- SciDAC software makes it very easy to get started, unlike general purpose super-computing centers where one might run into time consuming code porting.

Job submission

- All sites use some variant of PBS (Torque). As a result the usage is similar.
- PBS is also popular on several super-computing centers. Hence users are familiar with it.
- Detailed instructions with examples on how to prepare PBS scripts is provided on the web.
- QCDOC due to its nature requires special manipulations for submitting jobs. The needed steps are well documented on the web.
- Users need test queues for script testing and code development. All sites now provide such queues

Job monitoring

- Standard PBS commands allow users to check job and alter status.

- Both JLab and FNAL have nice web interfaces for monitoring the status of the cluster. Information related to job and queue status can be obtained through this interface.

Allocation Usage

- JLab: Very transparent web interface. User can check usage over specified past intervals
- JLab: Fair share scheduling is implemented
  - Users are allowed to consume time proportionally to yearly allocation. No cap to the total allowed time is imposed
    - Result: Lazy users lose their time while the diligent ones get more than their fair share.
- FNAL: No fair share scheduling. A user is allowed to consume up to its maximum allocated time.
  - If allocation is exceeded, user is allowed to submit to background queues in order to use free cycles.
- BNL: Hardware corresponding to users allocation is given
  - If user does not run hardware remains idle and allocation is used up

Data storage

- JLAB: Cache manager.
  - Large disks are automatically backed up on tapes. Files are deleted from disks if space is needed.
  - User can retrieve (srm commands) files from tape.
  - Files can be “pinned” if needed for a job completion
    - Effort is made to provide large enough disks to each project so that migration to tape does not become a burden to the users.
- FNAL: dcache plus large project associated disks.
  - User needs to manage his/her space.
- BNL: Large raid disks are provided. No archiving capability exists.
- Disk space is limited.... Users always need more, and if given what they ask will figure out a way to use it up!
- Usage of the systems is well documented on the web.

Problem Reporting

- BNL: CTS ticker system
  - Problem reports are logged.
  - A person is assigned to work on solving it. Reports are provided as work goes along
  - The complete history of a problem is available online
  - Personal experience: System worked very nicely. People responded quickly and efficiently in solving the problem
- JLAB and FNAL: Problems are reported by email (information exists on the web)
  - JLAB recently implemented an web based reporting system
  - For both JLAB and FNAL users reported (and personal experience) that the responce to problems is quick and effincient.

User Support

- In general users are happy with the support from the site managers
- Response is quick and helpful
- Users like the fact that support is offered by Lattice QCD theorists
- Special notice:
  - Don Holmgren, Amitoj Singh, and Jim Simone (FNAL)
  - Robert Edwards, Balint Joo, and Ying Chen (JLAB)
  - Bob Mawhinney, Chulwoo Jung, Stratos Eustathiadis (BNL)
- Advice is offered on both software development and issues related to job running, sometimes beyond the call of duty.

Problems

- Job efficiency: Generally reported job failure between 5% and 20% for all sites.
- This is consistent with behavior in other super-computing facilities
- On clusters NFS seems to be the weak point generating job failures and job stalling
- Bugs related to communication hardware divers are also creating issues
- QCDOC failure rate is a result of either hardware or software instabilities
- Instabilities of the tape backup systems are also reported
- Sometimes users need to spend time on data movement due to lack of adequate disk space. (Users may need to re-evaluate their computational strategy to avoid excessive disk usage)
- These problems don’t seem to affect much the science output of the systems

Conclusions

- The facilities are well managed and maintained.
- Users are generally happy and feel that they can do their research in a timely manner.
- Adequate documentation exists on the web for all aspects of the computational environment.
- Existence of SciDAC software seems very important in improving productivity.
- Reported problems seem to affect little the production output. However effort should be made to be addressed.